

Cellular Games: An Introduction

Lenore Levine

Department of Mathematics, University of Illinois,
273 Altgeld Hall, 1409 W. Green St., Urbana, IL 61801, USA

Abstract. A *cellular game* is a dynamical system in which cells, placed in some discrete structure, are regarded as playing a game with their immediate neighbors. Individual strategies may be either deterministic or stochastic. Strategy success is measured according to some universal and unchanging criterion. Successful strategies persist and spread; unsuccessful ones disappear.

In this paper, two cellular game models are formally defined, and are compared to cellular automata. Computer simulations of these models are presented.

Conditions providing maximal average cell success, on one- and two-dimensional lattices, are examined. It is shown that these conditions are not necessarily stable; and an example of such instability is analyzed. It is also shown that Nash equilibrium strategies are not necessarily stable.

1. Introduction

A *cellular game* is a dynamical system; that is, the variables it is composed of are regarded as changing over time. These variables or cells, arranged in a discrete structure such as a ring, are thought of as repeatedly playing a game with their neighbors. Most of this paper is concerned with one-dimensional cellular games, defined more formally as follows:

Definition 1.1. A one-dimensional cellular game *consists of*:

1. A one-dimensional discrete structure, uniform from the viewpoint of each site; that is, a ring or doubly infinite path.
2. A variable, or cell, at each site. The components of this variable may change at each discrete unit of time, or round. They consist of, at least:
 - (a) A move component, which can take on a finite number k of values.
 - (b) A strategy component, which determines what move a cell makes in a given round. A cell's strategy is based on past moves of it and its r nearest neighbors on each side. r is called the radius of the game. The number of past rounds considered is called the depth d of the strategy.

3. A fitness criterion, which does not change and is the same for each cell. This fitness criterion is usually local; that is, a cell's fitness in each round is based on its move and those of nearest neighbors within the radius of the game.
4. A mechanism for strategy selection, under which more fit strategies survive and spread. Strategy selection is usually nonlocal; that is, a more fit strategy may spread arbitrarily far in a fixed number of time units. An interval between strategy changes, which may be one or more rounds, is called a generation.

Thus, a cellular game can be considered a process in which cells make moves each round, based on their strategies, and strategies are updated in each generation, based on their fitness in preceding rounds.

Note that cellular game strategies, based on $kd(2r+1)$ possibilities, and fitness criteria, based on $k(2r+1)$ possibilities, are usually stored in the form of a table. Also note that n -dimensional cellular automata, with one cell for each n -tuple of integers or integers mod k , can be similarly defined.

One-dimensional cellular games are studied in [18], [2], [3], and [11]. Similar systems are discussed in [8], [9], and [10]; and games on a two-dimensional lattice in [13].

Cellular games satisfy a criterion for "artificial life" as discussed by Langton [5], that is, "There are no rules in the system that dictate global behavior. Any behavior at levels higher than the (individual cells) is, therefore, emergent."

Cellular games are a generalization and extension of well-known discrete dynamical systems called *cellular automata*. Cellular games were created largely because of questions arising from the observation of cellular automata. One-dimensional cellular automata are defined as follows.

Definition 1.2. A one-dimensional cellular automaton (CA) consists of:

- A one-dimensional discrete structure, uniform from the viewpoint of each site; that is, a ring or doubly infinite path.
- A variable, or cell, at each site, that can take on finitely many values or states. The cell's initial state may be specified as desired.
- A function that decides how each cell changes state from one generation, or discrete unit of time, to the next. This function, or cellular automaton rule, is always the same for each cell, and depends entirely on a cell's state and that of its r neighbors on each side in the past m generations. The parameter r is referred to as the radius of the cellular automaton, and m as its order. Cellular automaton rules are usually stored and described in the form of a table.

It can be shown that an m th-order CA is equivalent to a first-order CA having more states. This well-known proof, however, is dependent on the

locality of cellular automata—that is, on the fact that cells are directly affected only by their neighbors. For similar mathematical objects that are *not* local, such as cellular games, this proof cannot be used.

Theorem 1.3. *Any m th-order, k -state, r -radius cellular automaton is equivalent to a first-order cellular automaton with radius mr , that has mk states.*

Proof. Let C be an m th-order CA of k states and radius r . Let C' be a first-order, mk -state, mr -radius CA constructed as follows:

- Each m generations of C is considered a generation of C' .
- Each state of C' is considered a vector of m k -valued components.
- The j th component of cell c 's state is determined as follows:
 - If $j = 1$, by components $k - m + 1$ through k of the states of cells $c - r$ through $c + r$ in the previous generation.
 - If $1 < j < m$, by components $k - m + j$ through k of the states of cells $c - r$ through $c + r$ in the previous generation, and components 1 through $j - 1$ of these cells in the current generation.
 - If $j \geq m$, by components $j - m$ through $j - 1$ of the states of cells $c - r$ through $c + r$, in the current generation.

Automaton C' can simulate the actions of C ; it is trivially true that C' has mk states, and is first order. It thus suffices to show that C' has radius mr . To show this, however, it is enough to show that the j th component of cell c 's state is entirely determined by the states of cells $c - jr$ through $c + jr$ in the previous generation.

This statement is trivially true if $j = 1$. Now, assume this statement is true for $1, \dots, j - 1$. Consider the j th component of cell c 's state. The only thing it is dependent on, which did not influence previously considered components of this cell's state, is component $j - 1$ of cells $c - r$ through $c + r$ in this generation. And, by our induction hypothesis, these components are determined by the states of cells $c - r - (j - 1)r$ through $c + r + (j - 1)r$ in the previous generation. ■

Thus, if a cellular automaton of radius r operates on cells that can take k possible states, there are k^{2r+1} possible circumstances that need to be considered. The rule table, therefore, has k^{2r+1} entries; and there are $k^{k^{2r+1}}$ possible r -radius, k -state cellular automaton rules. An example of a cellular automaton rule is the two-state, radius-one rule whose evolution is illustrated below. In this rule, a cell can be in either state 0 or state 1. Any cell that in generation g is in state 1, and has both of its neighbors in state 1, stays in state 1 in generation $g + 1$. Otherwise, a cell is in state 0 in generation $g + 1$. This rule is Rule 128 according to Wolfram's [20] classification system of the 256 two-state, radius-one rules.

Generation 1:	1	0	1	1	1	1	0	1	0	1
Generation 2:	0	0	0	1	1	0	0	0	0	0
Generation 3:	0	0	0	0	0	0	0	0	0	0

Table 1: The action of Rule 128 on a circular ring of ten cells for three generations.

Definition 1.4. *A stochastic cellular automaton is as above, except that neighboring states do not determine the move made in the next generation, but rather the probability that a particular move will be made.*

Computer experiments on one-dimensional cellular automata are usually conducted with cells arranged in a ring. Cell states are indicated by colors; thus, k -state cellular automaton rules are often referred to as k -color rules. Initial conditions are displayed in a line on top of the screen, with each generation being displayed below the previous generation. In such experiments, initial conditions and rule table entries are often chosen with the aid of a pseudorandom number generator.

As a matter of fact, descriptions of computer experiments with cellular automata and other discrete dynamical systems often make reference, informally, to “random” initial conditions. This concept actually applies to mathematical models containing infinitely many variables, such as a one-dimensional cellular automaton with one cell for each integer. In such a case, “random,” “almost all,” or “normal” initial conditions refer to conditions such that all k^n n -tuples of k cell states are equally likely, for all n . Or, in other words, if the states of the cells are construed as decimal places of two real numbers, both numbers are normal to base k .

Such conditions cannot be duplicated exactly in the finite case, no matter how large the number of cells. However, conditions can be created that appear disordered and satisfy certain statistical tests of disorder. This is done with the aid of a pseudorandom number generator. Such initial conditions are often loosely referred to as “random.” Computer simulations of discrete dynamical systems often use such initial conditions as the most feasible indicator of likely behavior.

In such experiments, there are roughly three types of asymptotic behavior. First of all, all cells may become and remain one color, or change color periodically, with a small and easily observable period. Second, cells may display “chaotic” behavior; that is, cell color choice may appear to be disordered, or to result from some other simple stochastic algorithm. Third, cell color choice may be neither periodic nor chaotic, but appear to display organized complexity. That is, the cell evolution diagrams may look like biological structures, such as plants, or social structures, such as city maps. As a matter of fact, such diagrams are often quite aesthetically pleasing. These rule types are discussed in [21]; for more on the concept of “complexity,” as it applies to cellular automaton rules, see [19].

On a finite ring of cells, of course, all such evolution is eventually periodic. But, if cells can be in two states, and there are 640 cells, then there are 2^{640} possible ring states. Therefore the period of ring states could conceivably be quite high; and “chaotic” or “complex” rules do indeed seem to have very high periods.

Visual representations of cellular automata can exhibit a sophistication reminiscent of living structures. However, the number of k -state, r -radius cellular automaton rules is very large ($k^{k^{2r+1}}$) for all but the smallest k and r ; and “interesting” rules are not common and difficult to find. This leads to the question, therefore, of whether there is some way of “evolving” cellular automaton rules in a desired direction.

To this question, there are two possible avenues of approach. One is to select rules based on their global properties. That is, some computable measure of the desired characteristics is devised, and rules are chosen by their ability to meet this measure. Such selections are discussed in [15] and [12].

The other way is to select rules based on their local properties. That is, each cell uses a different rule; and there is some universal and unchanging criterion for rule success. This approach is more like the way living systems evolve, for the evolution of a planetary ecology is not due to constraints placed directly on the ecology. It is an emergent property of constraints placed on the individual organisms. For this reason, such models may potentially reveal not only the nature of “complex” rules, but also how their global properties emerge from local interactions.

An evolutionary model of this sort is equivalent to a cellular game; the only difference is the terminology. That is, a cell’s strategy can be regarded as the individual rule used by each cell; the depth of the strategy as the rule’s order; cell moves as states; and instead of referring to the smallest unit of time as a round, and a possibly larger unit as a generation, the smallest unit can, as with cellular automata, be referred to as a generation. The fitness criteria and evolutionary process stay the same.

A cellular game differs from a cellular automaton not only in the precise definitions used, but also in the philosophy under which these definitions were constructed. That is:

- Cellular automata are often regarded as physical models; for example, each cell may be seen as an individual atom. Thus, the rules by which each cell operates are the same. Cellular games, on the other hand, are seen as evolutionary models. Each cell uses an individual rule, or strategy, that can be thought of as the “genetic code” of the cell.
- Cellular automata are usually thought of as deterministic, beyond the initial generation, though stochastic CA’s have also been studied. Cellular games operate stochastically; that is, the evolutionary process under which strategies are modified is stochastic. Often the strategies themselves are too.
- Cellular automata are local; that is, a cell’s state is affected only by the states of its r nearest neighbors on each side in the previous generation.

In other words, cell information cannot travel more than r units per generation. This speed is often called “the speed of light.” Cellular games, on the other hand, typically use nonlocal strategy selection criteria. That is, a more fit strategy may propagate arbitrarily far in one generation.

- In Theorem 1.3, it is shown that m th-order CAs are behaviorally equivalent to first-order CAs with larger radius and more states. However, this proof does not work for cellular games with nonlocal selection criteria. Moreover, cellular games are often constructed with strategies that look more than one generation back.

It can be shown that if a cellular game has a local fitness criterion and local rule selection process, it is actually equivalent to a cellular automaton with a large number of states. This automaton, of course, will be stochastic if the game is stochastic.

Theorem 1.5. *Let G be a cellular game with a local fitness criterion and local rule selection process, which operates every g generations. Let all fitness measurements start over again after this process. Then G is equivalent to a cellular automaton G' with a much larger number of states.*

Proof. Let G' be constructed as follows. Let the state of a cell c in G' be a vector with the following components:

1. The state of c in G .
2. The individual rule used by c in G .
3. A g -valued counting variable, which starts out as 1 in the first generation, and thereafter corresponds to the current generation mod g .
4. A fitness variable, which corresponds to a cell's accumulated fitness.

Since these components enable G' to simulate the action of G , it suffices to show that G' is a cellular automaton. That is, each component must have only finitely many possible values, and be locally determined. Examining each component, in turn:

1. By definition of G , the first component has only finitely many values; and is determined by a cell's rule, and the states of it and its neighbors in preceding generations.
2. By Definition 1.1, even if stochastic rules are used, only finitely many are considered. Whether or not a cell keeps its rule, after g generations, is based on its own fitness, and the process of selecting new rules is assumed to be local.

3. The counting component can be in any one of g different states. The rule for its change is simple: If it is in state s in generation d , it is in state $(s + 1) \bmod g$ in generation $d + 1$. Note that to run G' as a simulation of G , this counting component must be initially set to the same value for all cells.
4. The fitness component is set to zero after every g generations, and can be incremented or decremented in only finitely many different ways. How it changes in each generation, for a given cell c , depends on the first components of cells $c - r$ through $c + r$. ■

Given this equivalence, why, then, is a cellular game so different from a cellular automaton? For one thing, a cellular game usually *does* use a nonlocal strategy selection process. For another, CA rule spaces are typically very large. Therefore, even if only systems with a local selection process are considered, the evolutionary paradigm of cellular games may still be valuable. It may be a practical method of selecting members of these spaces with interesting properties.

In this paper, two different models of cellular games are defined. The original Arthur-Packard-Rogers model is discussed first in Section 3. This model is quite extensive and uses many different parameters. The second, simplified model is more amenable to mathematical analysis. This model is discussed in Section 5.

Computer simulations of both models are presented. These simulations are similar to those of cellular automata, both in the way they are conducted and in the way they are displayed. That is, cell moves are indicated by colors. Strategies are usually not pictured, due to the large size of strategy spaces. Thus, a cell's move may also be referred to as its color. Initial moves of a finite ring of cells are displayed in a line on top of the screen, and each generation is displayed below the previous generation. Initial moves and strategies, as well as other stochastic choices during the course of the game, are implemented with the aid of a pseudorandom number generator.

Computer simulations of the first model display sophisticated behavior reminiscent of living systems, or "complicated" cellular automata. These behaviors, which include such phenomena as zone growth and "punctuated equilibria," are discussed and extensively illustrated in Section 4.

The second model admits only deterministic strategies of depth zero; that is, strategies of the form, "Do move m , without regard to previous rounds." Thus, in this model, moves and strategies can be considered equivalent. Though this model is simpler, there are still counterintuitive results associated with it. Even if only two strategies are allowed under this model, it is extremely difficult to predict which, if either, will be stable under invasion by the other. There are no simple algorithms for determining this.

For example, consider ring viability, discussed in Section 6. For finite rings this concept, Definition 6.1, refers to the average success of all cells in the ring. In this section, it is shown that under any local fitness criterion G , rings in which the cells have made periodic move sequences have the

highest possible viability. It is also shown that a similar result is false in the two-dimensional case.

Now, if cellular games did indeed always evolve towards highest ring viability, this would make their course relatively easy to predict. However, in Section 7, a two-strategy cellular game is presented, in which the best strategy for the ring as a whole—that is, the strategy that, if every cell follows it, maximizes ring viability—is not stable under invasion. This instability is illustrated by computer simulations. It is also proved mathematically. This is done by showing that if a small number of cells using the invading strategy are surrounded by large numbers that are not, the invading strategy tends to spread in the next generation. The reason for this is that the first strategy, though it does well against itself, does poorly against the second one.

On the other hand, a winning strategy may not necessarily be stable either. That is, strategy A may defeat strategy B, but still be unable to resist invasion by it. The reason, in this case, is that strategy B does so much better against itself. This result can also be demonstrated by computer simulations and proved, using the same method. These results are also in Section 7.

Finally, consider a situation in which, if its neighbors use strategy A, a cell has greatest success if it uses strategy A, too. It seems logical that, in this case, strategy A would indeed be stable. As a matter of fact, such a situation is called in game theory, a symmetric Nash *equilibrium*. However, it can be demonstrated by computer simulations, and also proved, that some symmetric Nash equilibrium strategies are *not* stable under invasion. The reason, in such cases, is that strategy B has somewhat less probability of surviving in a strategy A environment, but is very good at causing strategy A not to survive. Therefore strategy B is somewhat less likely to persist, but a lot more likely to spread. This result is also considered in Section 7.

Thus, the three theorems in Section 7 show how difficult it is to predict the course of cellular games, even under a very simple model. The counter-intuitive nature of the results obtained suggests the potential mathematical interest of this paradigm.

2. Game theory and cellular games

Success criteria in tabular form, or score tables, are extensively used in game theory. They describe the course of any game that can be exactly modeled, for which strategy success can be numerically described, and in which all strategies are based on finite, exact information. For example, consider the game of Scissors, Paper, Stone; that is, Scissors beats Paper, Paper beats Stone, and Stone beats Scissors. Suppose this game is played for one round, and the only possible strategies are deterministic. Then the table for this game is (if a win scores 1, tie at .5, and loss at 0):

Opponent	Scissors	Paper	Stone
Player			
Scissors	.5	1	0
Paper	0	.5	1
Stone	1	0	.5

The following definition is used in game theory.

Definition 2.1. *A mixed strategy is a stochastic strategy; that is, one under which, in some specified circumstances, more than one move has positive probability.*

A table can also be devised for mixed strategies, and for games of more than one round. For mixed strategies the table entry describes the expected success.

For example, suppose the game of Scissors, Paper, Stone is played for two rounds, and there are three possible strategies: Strategy A is to choose each move with probability $1/3$, Strategy B is to choose Stone for the first move, and the move chosen by the other player for the second, and Strategy C is always to choose Paper. Suppose, as above, the score is 1 per win per round and .5 per tie. Then the table for this game is:

Opponent	Strategy A	Strategy B	Strategy C
Player			
Strategy A	1	1	1
Strategy B	1	1	.5
Strategy C	1	1.5	1

Definition 2.2. *A table depicting strategy success as described above is called the normal form of a game.*

Normal form can be used, at least theoretically, to describe extremely sophisticated games. For example, if only a fixed finite number of moves is allowed, and strategies consider only the history of the current game, then there are only finitely many deterministic strategies for the game of chess. Hence normal form could, at least theoretically, be used to describe this game. Of course, there are so many possible chess strategies, this form cannot be used for practical purposes. For more on normal form, see [7].

Note that this form is ambiguous if mixed strategies are allowed. For example, consider the above table. Does it indicate the actual success levels of deterministic strategies, or the expected success levels of stochastic ones? It is not possible to tell without further information.

Such a normal form can also be used to describe three-player games. For example, this table describes a game in which there are two moves, you score 1 if you make the same move as both other players and 0 otherwise. This game is called the Join or Die game.

Your Move:	B		Your Move:	W	
Player 1:	B	W	Player 1:	B	W
Player 2:			Player 2:		
B	1	0	B	0	0
W	0	0	W	0	1

Now consider cellular games. If the success criterion, or score, is local—that is, if it is based entirely on a cell’s state and that of its neighbors—it can also be encoded as a table. As a matter of fact, any game table for $2r + 1$ players can be used as the score table for a cellular game of radius r . For example, the Join or Die process is a cellular game of radius 1, in which each cell plays the Join or Die game with its two nearest neighbors. The following table is used for this process:

Cell’s Move:	B		Cell’s Move:	W	
Right Neighbor:	B	W	Right Neighbor:	B	W
Left Neighbor:			Left Neighbor:		
B	1	0	B	0	0
W	0	0	W	0	1

However, cellular games differ from the situations most analyzed by game theorists, or the vernacular notion of a game, in the following ways:

- Each cell interacts with different neighbors, as determined by the discrete structure on which the cellular game is run. That is, cell 0’s score is based on its move, and those of cells 1 and -1 . Cell 1’s score is based on the moves of cells 0 and 2, not cells 0 and -1 .
- The “game” is considered to be played repeatedly, for many rounds. Thus, the main focus is on optimal move behavior in the long run, not for one round only.
- There is an explicit mechanism for determining how successful strategies thrive and spread. The cellular game is not completely described without this mechanism; no assumptions about asymptotic behavior can be made just on the basis of the score table.

3. The Arthur-Packard-Rogers model

The idea of cellular games was first developed by Packard and Arthur [14], and first written up by Rogers [18]. In this model, cells arranged in a ring play a game, such as the well-known Prisoner’s Dilemma, with each of their nearest neighbors. They play for a fixed number of rounds. At the end of these rounds, or of a generation, strategies may change. Successful strategies are most likely to spread and persist. (The Prisoner’s Dilemma is discussed in [16], [1], and Appendix B.)

The Arthur-Packard-Rogers model can be summarized as follows. Cells, arranged in a one-dimensional structure, play a game, such as the Prisoner's Dilemma, with their neighbors, for a predetermined number of rounds. The criteria for success in each round do not change, and are the same for each cell. Since the degree of success is based only on a cell's moves and those of its r nearest neighbors on each side, this criterion can be encoded in the form of a table.

For details of this model, see Appendix C. The terms used are described in Definition 1.1.

The strategies that govern cell move choices may be different for each cell, may be deterministic or stochastic, are based on past move history, and are stored in the form of a table. Strategies may have depth zero, one, or more.

At the end of these rounds—that is, at the end of a generation—the probability that a cell keeps its strategy in the next generation is proportional to the size of its reward variable, which measures its success in the game.

Definition 3.1. *Cell death: A cell is said to die if its strategy is so unsuccessful it must be replaced. The replacing strategy is usually derived from the strategies of other cells.*

Finally, if a cell dies at the end of a generation, the strategy chosen is some combination of the strategies of its nearest living neighbors. If it contains elements of both neighbors, crossover is said to occur.

Definition 3.2. *Crossover is the existence, in a new strategy, of elements from more than one “parent” strategy.*

Definition 3.3. *Those cells whose strategies contribute to a cell's new strategy are called its parents.*

There may also be a small probability of strategy table mutation.

Definition 3.4. *A mutation is said to occur when, after a strategy table entry has been chosen from a parent cell, it is arbitrarily changed.*

In computer simulations, this is often done with the aid of a pseudorandom number generator.

This model is not quite the same as the original one used in [18]. In that construction, strategy replacement was not governed by locality; that is, parent cells were the most successful in the ring. Thus, the progenitor of a cell's strategy was not particularly likely to be nearby.

In this model, however, parent cells are not necessarily the most successful cells in the ring. Instead, they are a cell's nearest living neighbors. Such a model is more comparable with living systems, because it bases system evolution more completely on local properties. It is also more easily generalizable to the infinite case, in which there is one cell for each integer. And it is only through such a model that one can see the evolution of different strategy zones.

4. Computer experiments

The Arthur-Packard-Rogers model has been simulated in computer experiments with the aid of a pseudorandom number generator. Cell moves are displayed on screen, in a form similar to the display of cellular automaton states. That is, initial moves for each generation are shown in a line on top of the screen; and moves for each round are shown below the preceding round. In experiments simulating the Prisoner's Dilemma, or variations, lighter areas indicate cooperative moves; dark areas, defecting moves. In particular, in the games illustrated below, all strategies are mixed, or stochastic. That is, there is a small probability that a move is made, other than the one called for by the strategy.

The experiment illustrated at the end of this paper, in Figures 1 through 14, simulates a variation of the Prisoner's Dilemma, the Stag Hunt. The Stag Hunt is modeled on the dilemma of a member of a pack of hunting animals, such as wolves or coyotes. If the whole pack hunts together, they can bring down a stag, which is the highest reward. If a member defects, it will be able to get a rabbit alone. If the other animals do not defect, they will have a smaller chance of bringing down a stag, but it may still be possible; but it is very unlikely that one animal can bring down a stag all by itself. Thus, the highest expected reward is for mutual cooperation; next highest, for defecting while the other members of the pack cooperate; next, for mutual defection, and fourth, for cooperating while the other members of the pack defect. See [16] for more information on the Stag Hunt, and Appendix A for a more technical discussion of the experiments.

These computer experiments fully suggest the mathematical interest of the subject. They reveal many kinds of thought-provoking behavior, such as:

- *Zone growth.* Strategies may not evolve in the same manner in all areas of the ring. Zones of cooperative, defecting, or other consistent behavior may arise and persist for generations.
- *Periodic structures.* Cells may alternate between cooperation and defection, or waves of cooperation may spread through some or all zones of the ring.
- *"Complexity."* Move patterns may display a sophistication reminiscent of living structures, or the patterns found in "complex" cellular automata.
- *Long transients.* Strategies predominant for hundreds of generations may ultimately disappear, and be replaced by completely different behavior.
- *"Punctuated equilibria."* Move behavior that appears to be stable for many generations may suddenly change very quickly—and then become stable again, for a long time.




Figure 1: Computer simulation of the Stag Hunt, a modified Arthur-Packard- Rogers cellular game model, with 75 cells and 150 generations per round. Program `cg2.pas`, random seed 824709136, generation 1. In this program, all initial strategies are depth 1, but strategies of depth up to 3 may be introduced as the system evolves.




Figure 2: The same program, parameters, and seed as in Figure 1, generation 27. Notice the rightward-moving waves of cooperative behavior in the right-hand part of the display. Here some zones exhibit cellular automaton-like triangular patterns.



Figure 3: Generation 139 of the run begun in Figure 1. Cellular automaton-like triangles predominate in this figure.



Figure 4: Generation 165. There are now leftward-moving waves of cooperative behavior in the middle of the display.



Figure 5: Generation 305.



Figure 6: Generation 483. Cellular automaton-like triangles appear again.




Figure 7: Generation 560. Move behavior does not appear to have changed much in many generations.



Figure 8: Generation 561. An all-cooperative zone appears. The next three figures show the rapid growth of this zone.




Figure 9: Generation 612.




Figure 10: Generation 658.




Figure 11: Generation 662. The all-cooperative zone has almost completely taken over the ring.



Figure 12: Generation 930. Large all-cooperative zones have predominated in the past several hundred generations. However, at this point, a perturbation in strategy—that is, an unexpected defect move—can set off many defect moves in other cells.



Figure 13: Generation 982. Recovery of an all-cooperative zone.




Figure 14: Generation 1262. At this point, perturbations do not set off much defecting behavior in other cells. That is, strategies are no longer “cooperate unless there are defectors in the neighborhood,” but “cooperate, whatever happens.”

5. The zero-depth model

These experiments well suggest the richness of behavior cellular games offer. The sophistication of patterns displayed provides ample justification for further study of this paradigm. But the Arthur-Packard-Rogers model does not lend itself well to mathematical analysis. Its implementation as a computer program is lengthy, and it contains many modifiable parameters. It is difficult to decide if any behavior exhibited is general, or just an artifact of the specific algorithms used.

To facilitate mathematical discussion of cellular game behavior, it is hence appropriate to simplify the model. Extensive study has been performed on such a model exhibiting the following simplifications:

- *Elimination of crossover.* The Arthur-Packard-Rogers model allows crossover. (Definition 3.2.)

In the simplified model, crossover is eliminated and each new strategy is an exact copy of one that already exists. A rationale for this simplification, in terms of living systems, is that one is considering the evolution of a specific gene, which spreads on an either-or basis. However, a particular gene may be significant only in the context of other factors. It may thus not be appropriate to consider this gene on its own. Note that computer experiments using genetic algorithms reinforce the significance of crossover (see [6]).

- *Elimination of mutation.* Another simplification is the elimination of mutation (Definition 3.4). That is, after the initial round, any strategy is new for a specific cell only, and is a copy of the strategy used by an existing cell. The elimination of mutation is quite likely to change the long-term behavior of the system, especially in the absence of crossover. For example, suppose strategy A is successful against all other strategies, including itself. If a ring of cells is originally free of strategy A, but mutation is allowed, strategy A will eventually take over the ring. If there is no mutation, the ring will stay free of it. However, the behavior of a cellular game that allows mutation may best be understood in terms of, and in comparison to, the behavior of the simpler system.
- *One round per generation.* That is, cell strategy may change after each round of play.
- *Elimination of mixed strategies.* Strategies are deterministic, not stochastic.
- *Elimination of depth.* The final simplification is the elimination of depth. That is, all strategies are executed without regard to past moves. Since there are no mixed strategies, the strategy, then, just becomes "do move m ," and the move variable can thus be eliminated from the description of the game.

The question of how depth and round restrictions affect process behavior is a subject for future research; however, these restrictions are not as severe as they seem. From game theory, we learn that all information about games with extremely sophisticated strategies can be conveyed in table form; that is, the “normal” form of a game. The only restriction is that strategies must take into account only a finite amount of information, for instance, the course of the game, but not anything before or beyond. As previously discussed, such tables can be used as the score table for a cellular game; in particular, for a zero-depth, one-round-per-generation cellular game.

As a matter of fact, cellular games of many rounds per generation, and with high-depth strategies, can be rewritten as zero-depth one-round games if all strategies take into account the current generation only.

Note that the Arthur-Packard-Rogers model does take into account moves in the previous generation. However, it could easily be modified not to do so, by providing table entries to use when there is limited information about previous rounds. For example, there could be an entry for the move used if nothing is known about previous moves.

Theorem 5.1. *Let G be a cellular game of radius r , with R rounds per generation, and strategies of depth d , except that all strategies take into account only moves in the current generation. Then the action of G can be exactly simulated by a cellular game G' of zero depth and one round per generation.*

Proof. It suffices to show that for every such game G there is a zero-depth, one-round cellular game G' , and a mapping f from strategies in G to strategies in G' , such that life probabilities correspond. Actions made after cell survival is decided can be the same in each case.

That is, suppose there are two rings of k cells each, $1 \leq k \leq \infty$. Let the first ring run G in generation g , and let each cell c use strategy S_c . Let the second ring run G' in that generation, and let each cell c' use strategy $f(S_c)$. Then the probability, at the beginning of g , that c survives into the next generation should be the same as the probability that c' does.

To show that such an f can be constructed, it suffices to show that the probability under G , at the beginning of a generation, that a cell will live through to the next generation is entirely dependent on its strategy and those of its $(R-1)r$ nearest neighbors on each side. For if this is true, a table can be constructed, giving the life probability for cell c if it and its neighbors follow strategies $S_{c-(R-1)r}, \dots, S_c, \dots, S_{c+(R-1)r}$. This table can be used to create a zero-depth, one-round cellular game with corresponding life probabilities.

Now life probabilities in G , at the end of a generation, are entirely dependent on that generation's move histories. Therefore, to show such strategy dependence, it is only necessary to show that the probability, at the beginning of g , that cell c will make move m in some subsequent generation q , is entirely dependent on c 's strategies and those of its $(q-1)r$ neighbors on each side.

This fact is trivially true in the first round of a generation. Since a cell has no information about past moves, the probability it makes move m is entirely dependent on its own strategy.

Now, suppose this fact is true for the first $q - 1$ rounds. In round q , the probability a cell makes move m is entirely dependent on its strategy, and the moves made by it and its r neighbors on each side in preceding rounds of this generation. Therefore, by the induction hypothesis, this probability at the beginning of a generation is entirely dependent on the strategies of the $(q-2)r$ neighbors of *these* cells, namely, cells $c - (q-1)r$ through $c + (q-1)r$. ■

We are thus left with the following model in which, associated with each cell c , in each generation g , are:

- A move/strategy variable $m_{c,g}$ from some finite alphabet Σ of k characters.
- A binary-valued life variable $L_{c,g}$. This variable can be set to either living, or not living.

After each generation, cell strategies change as follows:

- The probability that a cell's life variable is set to 1, so that it "lives" into the next generation, is determined by a universal and unchanging game matrix G . That probability is based on a cell's move/strategies and those of its r nearest neighbors on each side, in that generation.
- A live cell keeps its strategy in the next generation.
- A cell that does not live is given a new strategy in the next generation. This strategy is either that of its living nearest neighbor to the left, or to the right, with a 50% probability of each. If there are no living neighbors to either side, all possible strategies are equally likely.

This model lends itself easily to computer simulation, with the different strategies represented by different colors. Thus, in descriptions of this model, "move," "strategy," and "color" are equivalent. Such a simulation is presented at the end of this paper in Figure 15. In this simulation, a cell has probability 0.27 of living if it is the same color as both of its neighbors and 0.53 otherwise. Note that these probabilities do not necessarily sum to 1. Due to the shapes of the space-time zones produced, this process is called the Cloud Process.

Note that in the fifth generation, there are no more living white cells. Therefore, the nearest living neighbors of each cell are black, and every cell will become black in the sixth generation and stay that way from that point on.

We now discuss a theorem pertinent to this model; that is, a simple characterization of identity games. An identity game is a game in which, outside of certain pathological cases, no cell can change color. To avoid complications arising from these cases, the identity game is formally defined as follows:




Figure 15: Computer simulation of a one-round cellular game, the Cloud Process, on a ring of 640 cells. The table for this game is: $G(BBB) = G(WWW) = 0.27$, $G(BBW) = G(BWB) = G(BWW) = G(WBB) = G(WBW) = G(WWB) = 0.53$. Program `cloud.pas`, random seed 118950941. Initial conditions were chosen with the aid of a pseudorandom number generator.

Definition 5.2. *An identity game is a game in which, under at least some circumstances, cells have positive probability of living; and in which no cell can change strategy unless there are no living cells either to the left or right of it.*

The characterization alluded to is:

Theorem 5.3. *Under the zero-depth model, a cellular game is the identity game if and only if the probability that a cell stays alive is 1, provided its strategy is different from at least one of its neighbors.*

Proof. Suppose G is a zero-depth cellular game of radius r , with life probabilities fitting the above description. Suppose a cell has living neighbors on each side. Then either:

1. The cell is not the same color/strategy as both of its neighbors and it will stay alive.
2. The cell c is the same color as both of its neighbors, but has neighbors on both sides of different colors, the nearest ones being cells $c - r_1$ on the left and $c + r_2$ on the right. Then cells $c - r_1 + 1$ and $c + r_2 - 1$ are alive. Therefore, if c dies, c 's left parent will be cell $c - r_1 + 1$, or a cell closer to c ; and c 's right parent will be cell $c + r_2 - 1$, or a cell closer to c . Thus if c dies, both parents will be the same color as c .

On the other hand, suppose G is such that there is positive probability a cell c_1 of color a , next to a cell c_2 of color b , may not live. Let there be a configuration of cells giving positive life probability to the center cell. Thus, since life probabilities are determined locally, it is possible that there may be living cells on either side of c_1 . Let c_1 die, and let it have living neighbors on each side. If either of these neighbors is not the same color as c_1 , c_1 may change color; if both are, c_2 will change color. ■

Finally, if cellular games, as described above, are intended to model living systems, two questions arise. First, why is a new strategy a symmetric function of the strategies of both parents, instead of, for example, being more influenced by the strategy of the nearest parent?

One answer is that this process is intended to model sexual reproduction, in which a gene has an equal possibility of coming from each parent. Another is that if there is *positive* probability that each gene comes from each parent, the model may actually not behave very differently. Future research may settle this question.

The second question is, why nonlocality? That is, why not say that if a cell has no living neighbors near enough, it just stays dead in the succeeding generation? In this case, comparison with living ecosystems suggests that locality is more appropriate, but with a very large radius. That is, suppose there's a large die-off of organisms in one particular area. Then organisms from surrounding areas will rush in very fast, to fill the vacant area—but they cannot rush in infinitely far in one generation. Once again, future research may settle whether the simplified assumption of nonlocality actually creates different long-term behavior.

6. Ring and torus viability

The following theorem describes move behavior that results in optimal cell viability, for a whole ring of cells. It applies to all cellular games with a local life probability matrix; that is, all games in which the probability a cell "lives" into the next generation is determined by its moves, and those of its neighbors less than a given number r units away. It thus applies to the Arthur-Packard-Rogers model. However, it is here described in terms of the one-round model given in the previous section.

Definition 6.1. *The ring viability of a finite ring of cells C running a one-round game G , in generation g , is the average life probability of these cells in that generation after moves are made, but before the cells' life variable is actually set.*

Since C has finitely many cells, whose moves are from a specific finite alphabet, there is some combination of moves which will maximize this viability. For example, in a one-round version of the Stag Hunt game, ring viability will be maximized if all cells cooperate; and in some versions of

the Prisoner's Dilemma, ring viability will be maximized if cells alternate between cooperation and defection.

The result obtained is that this optimal arrangement is periodic. The following lemma is used in the proof:

Lemma 6.2. *Let G be a one-round cellular game of radius r , in which there are k possible moves from some finite alphabet Σ . Let t be any string in Σ^* . Let $L(t)$ be the average life probability of all cells in a ring of $|t|$ cells, such that the move of the i th cell is the i th character of t . Then, if b , w_1 , and w_2 are strings in Σ^* , $|b| \geq 2r$,*

$$L(bw_1bw_2) = \frac{L(bw_1) + L(bw_2)}{2} \quad (1)$$

Proof. Consider a ring of cells consecutively making the moves in bw_1bw_2 . Cells making moves from w_1 are more than r units away from cells making moves from w_2 . Therefore, these cells cannot influence each other's life probabilities. In the same way, b is large enough so the life probabilities of cells making moves in either copy of b can be influenced by cells making moves in w_1 or in w_2 , but not by both. Therefore the average life probability of all cells is the same as if they were considered to be in two different rings. ■

The main result follows.

Theorem 6.3. *Let G be a one-round cellular game as above. Then there is some nonnegative m and some sequence t of m moves, such that rings of nm cells, in which the moves of t are repeated n times, have the maximum ring viability for finite rings of any size.*

Proof. There are only a finite number of strings in Σ^* that either contain no more than $2r$ letters, or, when circularly arranged, contain no duplicate, nonoverlapping $2r$ -tuples. Let such strings be called "good"; and let t be any "good" string that maximizes $L(t)$. We wish to show that

$$L(t) = \max_{s \in \Sigma^*} L(s) \quad (2)$$

because, then, rings repeating the moves of t one or more times would have maximal viability.

Now, this is trivially true for s such that $|s| \leq 2r$, because all such s are good. Suppose it is true for all s such that $|s| < n$. We wish to show that it is true for s such that $|s| = n$.

If s is good, this is trivially true. Suppose s is not good. Then $s = bw_1bw_2$, $|w_1|, |w_2| \geq 0$, $|b| = 2r$. By Lemma 6.2,

$$L(w_1bw_2b) = \frac{L(bw_1) + L(bw_2)}{2} \quad (3)$$

And, by our induction hypothesis, we know that $L(bw_1) \leq L(t)$ and $L(bw_2) \leq L(t)$. ■

A corollary to this theorem is concerned with asymptotic viability of doubly infinite arrays of cells.

Definition 6.4. Let the asymptotic viability $L(I)$ of a doubly infinite array of cells I be measured as follows:

$$L(I) = \limsup_{n \rightarrow \infty} \frac{\sum_{i=-n}^n l(I_i)}{2n+1} \quad (4)$$

Corollary 6.5. Let I be a doubly infinite array of cells, and let $l(c)$ be the life probability of a cell c , given its move and those of its r neighbors on each side. If t is that finite string that maximizes $L(t)$, then $L(I) \leq L(t)$.

Proof. Consider what life probability cells n through $-n$ would have if they were arranged in a ring, instead of part of a doubly infinite lattice. The only cells that might have different life probability are cells $-n$ through $-n+r-1$ and n through $n-r+1$. And as n becomes larger, the contribution of these $2r$ cells to ring viability goes to 0. ■

In the two-dimensional case, however, a result similar to Theorem 6.3 is false. That is, there are two-dimensional cellular games, for which no finite torus can achieve maximal torus viability. This is not proven directly, but is a corollary of results about Wang tiles.

A Wang tile is a square tile with a specific color on each side. A set of Wang tiles is a finite number of such tiles, along with rules for which colors can match. For example, a red edge may be put next to a blue edge, but not a white edge. Such a set is said to tile the plane, if the entire plane can be covered by copies of tiles in the set, so that all edge matchings follow the rules. Robinson [17] showed that there are sets of Wang tiles that can tile the plane, but permit no periodic tiling.

Note that the set of tiles described by Robinson admits an “almost periodic” tiling. That is, for any positive integer N , the plane can be covered with these tiles periodically so that, under the given rules, the proportion of tiles having unmatching edges is less than $1/N$.

Now a two-dimensional cellular game can be made from a k -colored set of Wang tiles as follows. Let a cell be considered a tile, let there be k^4 possible moves, and let these moves be considered direct products of the colors of the Wang tiles. Let a cell’s life probability be increased by $1/4$ for every match of a component of its move, with the corresponding component of its neighbor’s move. For example, $1/4$ would be added to a cell’s life probability, if the left component of its move were compatible to the right component of its left neighbor’s move.

Suppose a cellular game were made in this manner from the set of tiles described by Robinson. Then no torus could have viability one, because otherwise there would be a periodic tiling of the plane using these tiles. However, there are periodic tilings of the plane for which only an arbitrarily small proportion of the tiles have unmatching edges. Therefore, since a periodic tiling of the plane can be considered a tiling of a torus, there are torus tilings having viability $1 - \epsilon$, for any $0 < \epsilon < 1$.

The comparison of cellular games and Wang tilings suggests other possibilities for future research on tilings. For example, instead of a Wang tiling in which two colors either match or not, one could consider a tiling in which two colors partially match. This would correspond to a cellular game in which more than two different levels of success were possible.

7. Strategy stability

In the preceding section, the concept of ring viability was discussed. That is, for each cellular game, there is some periodic combination of moves which maximizes average cell viability. One might assume that all cellular games would stabilize with cells exhibiting, or mostly exhibiting, such a combination of moves. If this assumption were true, questions about the long-term evolution of cellular games could be trivially resolved.

However, computer experiments simulating the Prisoner's Dilemma as a cellular game suggest that this is not necessarily the case. Specifically, a one-round cellular game is simulated in which each cell plays the Prisoner's Dilemma with each of its neighbors. Specifications are:

- *Radius.* The game is of radius one.
- *Strategies.* There are two strategies, or colors: "C," cooperate, or white; and "D," defect, or black.
- *Game Table.* The game life probability table is: $G(CDC) = 1$, $G(CDD) = G(DDC) = 7/10$, $G(CCC) = 6/10$, $G(DDD) = 4/10$, $G(CCD) = G(DCC) = 3/10$, $G(DCD) = 0$. (That is, $G(m_1m_2m_3)$ is a cell's probability of survival, if its right neighbor's move is m_1 , its own move is m_2 , and its left neighbor's move is m_3 .)

Under these circumstances, maximal ring viability is achieved by a ring of all-cooperating cells. And yet, computer experiments simulating this game do *not* show the mostly cooperative state to be stable. In the simulation depicted in Figure 16, if a small number of defecting cells are put in the middle of a large ring of cooperators, the defecting strategy quickly takes over the ring.

The answer to this puzzle is that, although defectors do badly against each other, they do extremely well against cooperators. Thus, if a small zone of defecting cells is placed in a large ring of cooperating cells, the zone between the leftmost and rightmost defecting cells tends to expand.

To address such questions more formally, we use the concept of a domain:

Definition 7.1. A domain is a contiguous row of same-colored cells.

We would like to examine what happens when a small defecting domain is placed between two very large cooperating domains. Is the number of defecting cells in the vicinity of that domain likely to go up or down? If it is




Figure 16: Computer simulation of a one-round Prisoner's Dilemma game on a ring of 600 cells. Initially, two defectors are placed side-by-side; all other cells are cooperators. (Black indicates defecting cells, and white, cooperating.) Program `prisoner.pas`, random seed 424479774. Note that the rate of expansion of the black domain appears roughly similar on each side, suggesting an informal estimate of the expected rate.

more likely to go up, we can reasonably say that cooperative behavior is not stable under invasion.

Conceivably, of course, *each* strategy could be unstable under invasion by the other; that is, there could be a tendency for large domains of each color to break up into smaller ones.

Let there be a doubly infinite lattice of cells, running the Prisoner's Dilemma game described above. Let B be a small, but greater than one-cell, black domain in this lattice, bordered, in generation 0, by two large white domains W_l and W_r . Let B_0 be the number of black cells in B in generation 0. Let B_1 equal B_0 plus the number of cells that were white in generation 0, and, in generation 1, have black strategies descended from the strategies of cells in B ; and minus the number of cells that were in B in generation 0, and are white in generation 1. Thus, B_1 is roughly the number of black cells in the vicinity of B in the next generation. Finally, let c_1 be the rightmost member of W_l , c_2 the leftmost member of B , c_3 the rightmost member of B , and c_4 the leftmost member of W_r , in generation 0.

Now, some terms used in this proof are defined.

Definition 7.2. *Let a black incursion be a situation in which a black cell c in D , becomes in the next generation the parent of newly black cells in W_l or W_r . If it becomes the parent of cells in both, let it be regarded as two incursions.*

Definition 7.3. Let the cell c , the parent of the newly black cells in the incursion, be called the parent of the incursion.

Definition 7.4. Let a white incursion, and its parent, be defined in a similar manner; that is, a situation in which a white cell becomes the parent of cells formerly in B .

Definition 7.5. Note that there cannot be a black incursion into W_l , unless c_1 dies (or into W_r , unless c_4 dies); and, similarly, there cannot be a white incursion into B , with parent in W_l , unless c_2 dies (or with parent in W_r , unless c_3 dies). Let such situations be called black (white) incursion possibilities.

We now show that as the size of the bordering white domain becomes arbitrarily large, the expected size of a black incursion into that domain (if possible, as explained above) should go to $5/6$.

Lemma 7.6. Let E_n be the expected size of a black incursion into a white domain W , given that there is a black incursion possibility with parent in B , and that $|W| = n$. Then, under G ,

$$\lim_{n \rightarrow \infty} E_n = \frac{5}{6}. \quad (5)$$

Proof. Suppose the nearest cell w , in W , to B to stay alive is such that there are k dead cells in W between w and B . Then cells in W between w and B have parents of both colors, and their probability of becoming black is thus $1/2$. Now, the probability of there being k such cells to die, under G , given the incursion possibility, is

$$G(CCC)[1 - G(CCC)]^{k-1} = \frac{3}{5} \left(\frac{2}{5}\right)^{k-1}.$$

(That is, each white cell with two white neighbors has probability $G(CCC) = 3/5$ of living.) Thus

$$\lim_{n \rightarrow \infty} E_n = \lim_{n \rightarrow \infty} \sum_{k=1}^n \binom{k}{2} \left(\frac{3}{5}\right) \left(\frac{2}{5}\right)^{k-1} = \sum_{k=1}^{\infty} \binom{k}{2} \left(\frac{3}{5}\right) \left(\frac{2}{5}\right)^{k-1} = \frac{5}{6}. \quad (6)$$

■

We also bound the expected size of a white incursion.

Lemma 7.7. Let E_m be (under G) the expected size of a white incursion into B , from a white domain W , given that there is a white incursion possibility with parent in W , and that $|B| = m$. Then $E_m < 5/4$.

Proof. Suppose the nearest cell b , in B , to W to stay alive is located so that there are k dead cells in B between b and B . Then cells in B between b and W have parents of both colors, and their probability of becoming white is thus $1/2$. Now, the probability of there being k such cells to die, under G , given the incursion possibility, is

$$G(DDD)[1 - G(DDD)]^{k-1} = \frac{2}{5} \left(\frac{3}{5}\right)^{k-1}.$$

(Since each black cell with two black neighbors has probability $2/5$ of living.) Thus

$$E_m = \sum_{k=1}^m \binom{k}{2} \left(\frac{2}{5}\right) \left(\frac{3}{5}\right)^{k-1} < \sum_{k=1}^{\infty} \binom{k}{2} \left(\frac{2}{5}\right) \left(\frac{3}{5}\right)^{k-1} = \frac{5}{4} \quad (7)$$

■

The main theorem follows.

Theorem 7.8. *Let B be a small black domain on a doubly infinite lattice, on which the Prisoner's Dilemma game G is run. Let all variables be as described above. Then, if $B_0 \geq 2$, and W_l and W_r are large enough, the expected value of $B_1 - B_0$, which is roughly the expected change in the number of black cells in the vicinity of W , is positive.*

Proof. We examine eight cases, depending on the life of c_1 , c_2 , c_3 , and c_4 . Note that c_1 and c_4 have probability $G(CCD) = G(DCC) = 3/10$ of living; and c_2 and c_4 , $G(CDD) = G(DDC) = 7/10$.

1. All four cells live. Then $B_1 - B_0 = 0$.
2. c_1, c_2, c_3 live, c_4 does not (or the reflection of this case). The probability of this is $2(3/10)(7/10)^3$. There is one black incursion possibility (with c_3 as the parent) of expected size that goes to $5/6$, as the neighboring domain becomes arbitrarily large.
3. c_1, c_2 live, c_3 dies, c_4 lives (or the reflection). The probability of this is $2(3/10)(7/10)(3/10)^2$. There is one white incursion possibility (with c_4 as the parent) of expected size less than $5/4$.
4. c_1, c_2 live, c_3, c_4 die (or the reflection). The probability of this is $2(3/10)(7/10)(3/10)(7/10)$. There is one black incursion possibility (with c_2 or a cell between c_2 and c_3 as the parent) of expected asymptotic size $5/6$; and there may be one white incursion possibility (with a cell to the right of c_4 as the parent), of expected size less than $5/4$.
5. c_1 dies, c_2 lives, c_3 lives, c_4 dies. This case has probability $(7/10)^4$. There are two black incursion possibilities (with c_2 and c_3 as the parents) of expected asymptotic size $5/6$ each.

6. c_1 dies, c_2 lives, c_3 dies, c_4 lives (or the reflection). The probability of this is $2(7/10)^2(3/10)^2$. There is one black incursion possibility (with parent c_2) of expected asymptotic size $5/6$; and one white incursion possibility (with parent c_4) of expected size less than $5/4$.
7. c_1 dies, c_2 lives, c_3 and c_4 die (or the reflection). The probability of this is $2(7/10)^2(3/10)(7/10)$. There is one black incursion possibility (with parent c_2) of asymptotic size $5/6$; and there may be one white incursion possibility (with parent to the right of c_4) of expected size less than $5/4$.
8. c_2 and c_3 both die. The probability of this is $(3/10)^2$. There may not be a black incursion if every cell in D dies. There are at most two white incursion possibilities of expected size less than $5/4$ each.

Thus, if $|B| \geq 2$, and W_l and W_r are large enough, under all cases the expected value of $B_1 - B_0$ must be at least

$$\begin{aligned}
 & 2 \left(\frac{3}{10} \right) \left(\frac{7}{10} \right)^3 \left(\frac{5}{6} \right) - 2 \left(\frac{7}{10} \right) \left(\frac{3}{10} \right)^3 \left(\frac{5}{4} \right) + 2 \left(\frac{7}{10} \right)^2 \left(\frac{3}{10} \right)^2 \left(\frac{5}{6} - \frac{5}{4} \right) \\
 & + \left(\frac{7}{10} \right)^4 2 \left(\frac{5}{6} \right) + 2 \left(\frac{7}{10} \right)^2 \left(\frac{3}{10} \right)^2 \left(\frac{5}{6} - \frac{5}{4} \right) \\
 & + 2 \left(\frac{7}{10} \right)^3 \left(\frac{3}{10} \right) \left(\frac{5}{6} - \frac{5}{4} \right) - \left(\frac{3}{10} \right)^2 2 \left(\frac{5}{4} \right) = \frac{841}{6000}. \blacksquare
 \end{aligned}$$

However, it is not always the case that, in a two-strategy system, the “dominant” strategy will prevail. One strategy may lose against another, but do so well against itself that its use tends to expand. This happens in zero-depth versions of the previously discussed Stag Hunt, a game similar to the Prisoner’s Dilemma, except that successful cooperation is more profitable than exploitation. If computer experiments (Figure 17) simulate this game, giving a high enough premium for mutual cooperation, then cooperative behavior does tend to prevail. Specifically, the game has the same radius and number of moves as the Prisoner’s Dilemma game described above. Its table is: $G(CDC) = 10/16$, $G(CDD) = G(DDC) = 7/16$, $G(CCC) = 1$, $G(DDD) = 4/16$, $G(CCD) = G(DCC) = 8/16$, $G(DCD) = 0$.

It is possible, using the same techniques as above, to show that black domains are unstable in this game.

Theorem 7.9. *Let W be a small white domain on a doubly infinite lattice, on which the Stag Hunt game as described above is run. Let B_l and B_r be its neighbors, and W_0 be its size in generation 0. Let W_1 equal: W_0 plus the number of cells that were black in generation 0, and, in generation 1, have white strategies descended from the strategies of cells in W ; and minus the number of cells that were in W in generation 0, and are black in generation 1. If $W_0 \geq 2$, and B_l and B_r are large enough, then the expected value of $W_1 - W_0$, roughly the expected change in the number of white cells in the vicinity of W , is positive.*




Figure 17: Computer simulation of a zero-depth, one round per generation Stag Hunt game on a ring of 600 cells. Initially, four cooperators are placed contiguously; all other cells are defectors. Program `stag.pas`, random seed 941165838. Note that, in this case, the rate of expansion of the white domain appears to vary considerably.

Proof. The same calculations as described above are carried out, except that white and black are exchanged, and the probabilities of the Stag Hunt game are used. The asymptotic expected size of a white incursion, given the possibility of such, turns out to be 2. The expected size of a black incursion, given the possibility of such, turns out to be at most $1/2$ (since cells that are white and bordered on both sides by white neighbors cannot die). The asymptotic expected change in the number of white cells in the vicinity of W turns out to be at least $223/256$. ■

Nash equilibria of cellular games have also been analyzed [2].

Definition 7.10. *In the context of a cellular game, a symmetric Nash equilibrium (SNE) arises if, when a cell's r nearest neighbors on each side use strategy s , its best response is also to use s .*

For example, in the Stag Hunt game described above, both unilateral cooperation and defection give rise to such equilibria. That is, if a cell's neighbors always cooperate (defect), a cell is best off cooperating (defecting) too.

As with ring viability, it is easy to assume that Nash equilibria determine the course of a game; that is, a strategy giving rise to a symmetric Nash equilibrium is stable under invasion by differently behaving strategies. However, while the study of Nash equilibria is a promising avenue to understanding cellular games, such an automatic assumption is not necessarily the case.



Figure 18: Computer simulation of a zero-depth, one round per generation Stag Hunt game on a ring of 600 cells. Initially, four defectors are placed contiguously; all other cells are cooperators. Program `stag2.pas`, random seed 90049811.

For example, in the Stag Hunt, unilateral cooperation gives rise to a SNE. However, in some versions of this game, cooperating domains are unstable because isolated defecting cells are quite likely to kill off their neighbors despite the fact they tend to die off themselves. Thus, they tend to have more descendants than their neighbors.

The parameters used in this version of the Stag Hunt are not exactly the same as above. They are: $G(CDC) = 16/18$, $G(CDD) = G(DDC) = 15/18$, $G(CCC) = 1$, $G(DDD) = 14/18$, $G(CCD) = G(DCC) = 9/18$, $G(DCD) = 0$.

Computer experiments simulating this process (Figure 18) do indeed show that white domains are unstable. This result can also be proved using the same techniques as above.

Theorem 7.11. *Let B be a small black domain on a doubly infinite lattice, on which the second Stag Hunt game as described above is run. Let W_l and W_r be its neighbors, and let B_0 be its size in generation 0. Let B_1 equal: B_0 plus the number of cells that were white in generation 0, and, in generation 1, have black strategies descended from the strategies of cells in B ; and minus the number of cells that were in B in generation 0, and are white in generation 1. If $B_0 \geq 2$, and W_l and W_r are large enough, then the expected value of $B_1 - B_0$, roughly the expected change in the number of black cells in the vicinity of B , is positive.*

Proof. The same calculations as described for the Prisoner's Dilemma case are carried out, except that the probabilities of the second Stag Hunt game

are used. The asymptotic expected size of a black incursion, given the possibility of such, turns out to be $1/2$, since cells that are white and bordered on both sides by white neighbors cannot die. The expected size of a white incursion, given the possibility of such, turns out to be at most $9/14$. The asymptotic expected change in the number of black cells in the vicinity of B turns out to be at least $311/1008$. ■

Thus we see that cellular game behavior is difficult to anticipate. These systems reflect the richness of living ecologies, in which how well a species survives is determined by how well the organisms of that species compete with others, how well they cooperate among themselves, and how many descendants they have. No one factor automatically decides the issue.

8. Acknowledgments

I would like to thank Julian Palmore for extensive suggestions and discussion, and Norman Packard for introducing me to this new and challenging area. I would also like to thank Larry Dornhoff for help with the computers.

Appendix A. Computer experiments

All computer experiments were done in Turbo Pascal, Version 4.0, using the built-in pseudorandom number generator. Source code is available from levine@symcom.math.uiuc.edu.

The program simulating the modified Arthur-Packard-Rogers model, with Stag Hunt parameters, is `cg2.pas`. Note that in this program all strategies are mixed; that is, there is a small probability of actions other than those called for by the pure strategy.

The implementations of one-round models are as follows: The Cloud Process, `cloud.pas`, the Prisoner's Dilemma, `prisoner.pas`, the Stag Hunt (first version), `stag.pas`, and the Stag Hunt (second version), `stag2.pas`.

Appendix B. The Prisoner's Dilemma

The Prisoner's Dilemma is a two-person game in which two types of moves are possible: cooperate and defect. This game models the options of two prisoners held in separate cells for the same crime, who are being pressured to confess to that crime. If both prisoners keep silent—that is, they cooperate with each other—they will both get a small sentence for a lesser crime. If they both talk—that is, they both defect—they both get the standard sentence. If one talks and the other does not, the one that kept silent gets a very severe sentence and the other goes free. Thus, Prisoner's Dilemma is a game in which a player's reward for defecting, while the other player cooperates, is highest. Next highest is the reward for mutual cooperation; then, the reward for mutual defection. Lowest of all is the reward for cooperating while the other player defects. The Prisoner's Dilemma can also be generalized to three-person games.

Appendix C. The Arthur-Packard-Rogers model

The computer experiments presented in section 3 use a model very similar to the one described in [18]. That is, there exists a circular ring, or doubly infinite lattice, of cells C . Associated with each cell c , in each round i of each generation g , are:

- A move variable $m_{c,i,g}$ from some finite alphabet Σ of k characters.
- A strategy variable $S_{c,g}$. This is a table, in which entries are from Σ . If strategies are of depth d and radius r (that is, moves of a cell's r nearest neighbors, up to d rounds back, are taken into account), then this table contains $k^{d(2r+1)}$ entries. Hence there are $k^{d(2r+1)}$ possible strategies. Note that strategies do not change in a generation, but they do take into account rounds in previous generations. In computer experiments, move and strategy variables are initialized with the aid of a pseudorandom number generator.

A finite number of *mixed*, that is, stochastic, strategies may also be implemented; that is, strategies in which, given at least one game history, there is positive probability of a cell making two different moves. For example, a mixed strategy for Prisoner's Dilemma would be to cooperate 95% of the time, and defect the other 5%. If a given game allows k moves, and k' mixed courses of action, there are $(k+k')^{k^{d(2r+1)}}$ possible strategies. Again, mixed strategies, and all other stochastic actions, are implemented with the aid of a pseudorandom number generator.

- A reward, or payoff variable $W_{c,i,g}$. This variable starts out at 0 in the first round of each generation: its change in each round measures the cell's success in that round.

Changes to the reward variable are determined by a matrix G that defines the game and does not change during its course. That is, if a cellular game has radius r , and $i > 1$,

$$W_{c,i,g} = W_{c,i-1,g} + G(m_{c-r,i,g}, \dots, m_{c,i,g}, \dots, m_{c+r,i,g}) \quad (8)$$

An example of a game matrix is the following table for a Prisoner's Dilemma game. If "D" is defect, and "C" is cooperate: $G(CDC) = 100$, $G(CDD) = G(DDC) = 70$, $G(CCC) = 60$, $G(DDD) = 40$, $G(DCC) = G(CCD) = 30$, $G(CDC) = 0$. For this game, $k = 2$ (that is, there are two possible moves, cooperate or defect); and $r = 1$ (only the moves of a cell's nearest neighbors affect its reward variable).

In the Arthur-Packard-Rogers model, a fixed number of rounds R (e.g., 150 rounds), is regarded as constituting a generation. After each generation, cell strategies change, as follows:

- A cell's probability of "living" into the next generation is an increasing function of the size of its reward variable. Usually the reward matrix contains only positive entries, and life probability is proportional to the size of a cell's reward variable.

- A live cell keeps its strategy in the next generation.
- A cell that does not live is given a new strategy in the next generation. This strategy is chosen as follows:
- New entries in the strategy table are taken from corresponding entries from either one of the two *parent* cells (the cell's nearest living neighbors on each side). The new strategy table can contain elements from both parent cells (crossover, Definition 3.2) or one parent (no crossover). The exact details of how such a selection is carried out is part of the *genetic algorithm* used in the program. For a discussion of genetic algorithms, see [4]. Note, however, that all such algorithms are symmetric between the left and right parent; and that if a cell has no living neighbors on either side, all strategy possibilities are equally likely.
- After the basic new strategy is chosen, each table entry is subject to *mutation* (Definition 3.4). That is, there is a small probability it may change.

References

- [1] R. Axelrod, *The Evolution of Cooperation* (New York, Basic Books, 1984).
- [2] R. Cowan and J. H. Miller, "Life on a Lattice: The Nature of Equilibria in Spatially Overlapping Games" (unpublished).
- [3] R. Cowan and J. H. Miller, "Economic Life on a Lattice: Some Game Theoretic Results," Santa Fe Institute Working Paper 90-10 (1990).
- [4] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Reading, MA: Addison-Wesley, 1989).
- [5] C. G. Langton, *Artificial Life* (Reading, MA: Addison-Wesley, 1989).
- [6] S. Levy, *Artificial Life* (New York: Pantheon, 1992).
- [7] R. D. Luce and H. Raiffa, *Games and Decisions* (New York: Wiley, 1957).
- [8] K. Matsuo and N. Adachi, "Metastable Antagonistic Equilibrium and Stable Cooperative Equilibrium in Distributed Prisoner's Dilemma Game," *Proceedings of the International Symposium on Systems Research, Informatics and Cybernetics*, Baden-Baden (1989).
- [9] K. Matsuo and N. Adachi, "How to Attain to Cooperative Society in Game World: The Choice of Selection Rules," The International Institute for Advanced Study of Social Information Science, Fujitsu Laboratories, Ltd., Japan.
- [10] K. Matsuo and N. Adachi, "Ecological Dynamics under Different Selection Rules in Distributed and Iterated Prisoner's Dilemma Game," The International Institute for Advanced Study of Social Information Science, Fujitsu Laboratories, Ltd., Japan.

- [11] J. H. Miller, "The Evolution of Automata in the Repeated Prisoner's Dilemma" (University of Michigan, 1988).
- [12] M. Mitchell, P. T. Hraber, and J. P. Crutchfield, "Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations," *Complex Systems*, **7** (1993) 89–130.
- [13] M. A. Nowak and R. M. May, "Evolutionary Games and Spatial Chaos," *Nature*, **359** (1992) 826.
- [14] N. Packard, personal communication (1993).
- [15] N. Packard, "Adaptation Toward the Edge of Chaos," pages 293–301 in *Dynamic Patterns in Complex Systems* (Singapore: World Scientific, 1988).
- [16] W. Poundstone, *The Prisoner's Dilemma* (New York: Doubleday, 1993).
- [17] R. M. Robinson, "Undecidability and Nonperiodicity for Tilings of the Plane," *Inventiones Mathematicae*, **12** (1971) 177–209.
- [18] K. C. Rogers, "Cellular Automata Simulations Exhibiting an Evolutionary Increase in Complexity," Master's Thesis, University of Illinois, Department of Electrical Engineering (1990).
- [19] M. M. Waldrop, *Complexity* (New York: Simon and Schuster, 1992).
- [20] S. Wolfram, "Statistical Mechanics of Cellular Automata," *Review of Modern Physics* **55** (1983) 601–644.
- [21] S. Wolfram, "Universality and Complexity in Cellular Automata," *Physica D*, **10** (1984) 1–35.